

Remote Replication of SAP® Systems on Microsoft® SQL Server® 2008 Environments Using the Hitachi Virtual Storage Platform

Implementation Guide

By Archana Kuppuswamy

July 2011



Feedback

Hitachi Data Systems welcomes your feedback. Please share your thoughts by sending an email message to SolutionLab@hds.com. Be sure to include the title of this white paper in your email message.

Table of Contents

Tested Solution Components	4
Hardware Components	6
Software Components.....	8
Solution Implementation	11
Configure the Storage Area Network	11
Configure Servers and the Operating Environment	13
Configure Storage for SAP ERP server at remote site.	14
Install and Configure SAP ERP on Server at Remote Site	15
Prepare Storage Systems for Hitachi Universal Replicator Software	15
Configure Hitachi Command Control Interface (CCI) for Replication.....	20
Implement Remote Replication for SAP ERP Server	23
Appendix A—Hitachi Universal Replicator Configuration Files	26
HORCM Configuration File for SAP ERP Server at Local Site	27
HORCM Configuration File for SAP ERP Server at Remote Site	27
Appendix B—Assigning Authorizations for SQL Server Logon After Replication	28

Remote Replication of SAP® Systems on Microsoft® SQL Server® 2008 Environments Using the Hitachi Virtual Storage Platform

Implementation Guide

Most SAP customers rely on their IT environments to be available 24 hours a day, 7 days a week. Given their critical nature, most SAP deployments need a secondary copy of the production landscape's data. Rapid data recovery is critical for SAP operations with aggressive recovery time objectives (RTO) and recovery point objectives (RPO).

One of the most powerful capabilities of an advanced storage infrastructure is the ability to replicate data from the production instance at a local site to a secondary instance at the remote site. Data replication is an important part of an effective recovery policy from catastrophic system failure. Remote data replication offers a low risk of data loss.

Remote data replication can be one of the most complex storage administration tasks. Introducing any process into a production SAP environment requires detailed planning and testing to ensure that the replication process does not impact production. Although data replication can be performed at the server level, this function is performed more effectively within the storage infrastructure.

Hitachi Data Systems takes a holistic approach to SAP environments. The key to that approach is the Hitachi Virtual Storage Platform. The advanced functionalities of the Hitachi Virtual Storage Platform fulfill the complex and demanding requirements of an SAP implementation, including data replication.

All SAP solutions include an embedded database. Microsoft Windows Server and Microsoft SQL Server support secure, reliable, and scalable SAP environments.

This solution does the following:

- Replicate a SAP production ERP instance on Microsoft Windows Server 2008 R2 and Microsoft SQL Server 2008 R2
- Prepare the replicated data for access during a disaster or a planned failover by a server that is an exact copy of the primary local site server at a remote site.

This installation guide provides step-by-step instructions for preparing the environment for replication, configuring replication, and implementing replication in a SAP environment.

These instructions are written for IT professionals who are charged with the storage, deployment, or administration of SAP systems on the Hitachi Virtual Storage Platform and Hitachi Compute Blade 2000 using Microsoft Windows Server 2008 R2 and Microsoft SQL Server 2008 R2. It assumes the following:

- Familiarity with storage area network (SAN) technologies, tools, operating practices
- Basic knowledge of Hitachi storage management software, including Hitachi Command Suite and Hitachi Dynamic Link Manager Advanced

Tested Solution Components

This solution includes SAP clients that access the SAP application infrastructure at the local site. The SAP application layer stores data in and accesses data from a Hitachi Virtual Storage Platform family storage system at the local site. These SAP clients also can access the SAP application infrastructure at a remote site, which stores and accesses data from the Hitachi Virtual Storage Platform storage system at the remote site. Specifically, this solution uses the following:

- The platform for building the SAP environment in a local and remote location:
 - Hitachi Virtual Storage Platform
 - Hitachi Compute Blade 2000—The Hitachi Compute Blade 2000 chassis has the following on three server blades:
 - SAP Solution Manager
 - SAP ERP
 - SAP ERP off-host instance
- Hitachi Universal Replicator

The SAP client tier consists of Microsoft Windows servers with the SAP graphical user interface installed. The clients connect to the SAP system in the SAP application tier using a Cisco Catalyst 4500 switch at the local and remote sites. Hitachi Storage Navigator is installed on a management server in the client tier.

The SAP application and database tier consists of the SAP ERP and the SAP Solution Manager systems installed on Microsoft Windows servers on a central system at the local and remote site. This means that the SAP application and SAP database are installed on a single server.

All the SAP application servers are connected to the Hitachi Virtual Storage Platform using two Brocade 5300 switches.

Figure 1 illustrates the environment used to test this solution in the Hitachi Data Systems lab.

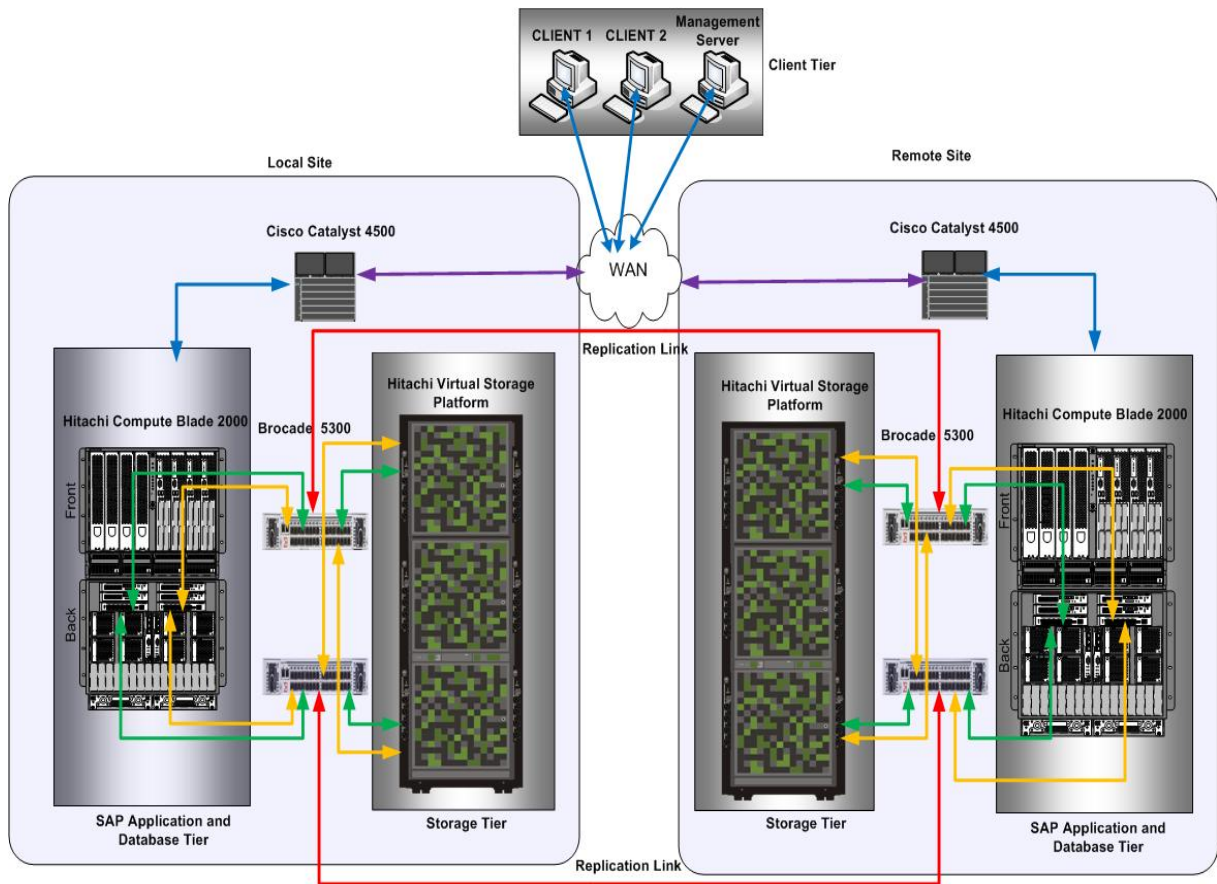


Figure 1

“Hardware Components” and “Software Components” describe the components required to deploy this solution.

Hardware Components

Table 1 describes the hardware used to deploy this solution in the Hitachi Data Systems lab.

Table 1. Tested Deployment Hardware

<i>Hardware</i>	<i>Quantity</i>	<i>Configuration</i>	<i>Role</i>
Hitachi Virtual Storage Platform storage system	2	8 Fibre Channel ports used 2 pair of front-end directors 2 pair of back-end directors 64 x 300GB 10K RPM SAS drives 64GB cache	Primary and Remote storage
Hitachi Compute Blade 2000 chassis	2	8-blade chassis 2 Fibre Channel switch modules 8 x 1Gb/sec network ports 2 x management modules 8 x cooling fan modules 4 x power supply modules	Blade Chassis
Hitachi Compute Blade 2000 server blade	1	2 x 8 Core processor with 64GB of memory	SAP Solution Manager server at local site
Hitachi Compute Blade 2000 server blade	1	2 x 8 Core processor with 64GB of memory	SAP ERP server at local site
Hitachi Compute Blade 2000 server blade	1	2 x 8 Core processor with 64GB of memory	SAP Solution Manager server at remote site
Hitachi Compute Blade 2000 server blade	1	2 x 8 Core processor with 64GB of memory	SAP ERP server at remote site
Brocade 5300 SAN Fibre Channel switch	2	FOS 5.3.1a 14 4Gb/sec Fibre Channel ports	SAN switch
Server	2	Intel Pentium 3.0GHz processor 1GB memory	SAP clients
Server	1	Intel Pentium 3.0GHz processor 1GB memory	Management server

Hitachi Virtual Storage Platform

The Hitachi Virtual Storage Platform is a 3D scaling storage platform. With the unique ability to scale up, scale out, and scale deep at the same time in a single storage system, the Virtual Storage Platform flexibly adapts for performance, capacity, connectivity, and virtualization.

- **Scale Up**— Increase virtual server consolidation, improve utilization of resources, and reduce costs.
- **Scale Out**— Combine multiple chassis into a single logical system with shared resources.
- **Scale Deep**— Extend the advanced functions of the Virtual Storage Platform to external multivendor storage.

For more information, see the [Hitachi Virtual Storage Platform](#) on the Hitachi Data Systems website.

Hitachi Compute Blade 2000

The Hitachi Compute Blade 2000 is an enterprise-class blade server platform. It features the following:

- A balanced system architecture that eliminates bottlenecks in performance and throughput
- Embedded logical partition virtualization
- Configuration flexibility
- Eco-friendly power-saving capabilities
- Fast server failure recovery using a N+1 cold standby design that allows replacing failed servers within minutes

Hitachi embeds logical partitioning virtualization in the firmware of the Hitachi Compute Blade 2000 server blades. This proven, mainframe-class technology combines Hitachi's logical partitioning expertise with Intel VT technologies to improve performance, reliability, and security. Embedded logical partition virtualization does not degrade application performance and does not require the purchase and installation of additional components.

SAP Servers

Servers hosting SAP Solution Manager and SAP ERP server must meet the specification requirements for the SAP roles they are hosting. For more information about server requirements for SAP ERP, SAP Solution Manager and SAP clients, see the following SAP Notes and websites, which are available to licensed customers from SAP's web site:

- SAP Note 26147, SAP GUI Resources: Hardware and Software
- SAP Note 901070, Resource Requirements for SAPECC6
- <http://service.sap.com/sizing>

SAP offers numerous architectural designs to maximize performance and suit the individual needs of SAP customers. Designing a SAP landscape requires proper planning by you and the solution implementation partner. For more information, consult with your implementation partner after reading the proper planning guides for SAP ERP 6.0 at the [SAP Service Marketplace](#), available to licensed SAP customers:

- **First Steps Document**—Helps you get started with the installation of the SAP ERP 6.0 instance.
- **Master Guide**—This contains information about the sequence and the components available for SAP ERP 6.0. It also contains links to all of the SAP ERP 6.0 documentation from SAP.

Software Components

Table 2 lists the software used to test this solution in the Hitachi Data Systems lab

Table 2. Tested Deployment Software

<i>Software</i>	<i>Version</i>
Hitachi Storage Navigator	Microcode dependent
Hitachi Command Suite	7.0
Hitachi Command Control Interface (CCI)	01-24
Hitachi Universal Replicator	Licensed feature available with Hitachi Virtual Storage Platform
Hitachi Dynamic Provisioning Software	Licensed feature available with Hitachi Virtual Storage Platform
Hitachi Dynamic Link Manager Advanced	6.0.1.0.804
SAP ERP	ERP 6.0 Support Release 1 (SR1) with Enhancement Pack 4 IDES, Upgraded SAP Kernel
SAP Solution Manager	7.0 Enhancement Package 1 on SP 25
SAP GUI	7.2
Microsoft SQL Server	2008 R2
Microsoft Windows	2008 R2

Hitachi Storage Navigator

Hitachi Storage Navigator is the integrated, web-accessible graphical management interface for the Hitachi Virtual Storage Platform firmware and software features. Use it to take advantage of all of the features of the Virtual Storage Platform.

Use Hitachi Storage Navigator to do the following:

- Map security levels for SAN ports and virtual ports
- Inter-system path mapping
- Logical unit (LDEV) creation and expansion
- Online volume migrations
- Configure and manage Hitachi replication products
- Perform online microcode updates and other system maintenance functions
- Integrate SNMP with enterprise management systems

Hitachi Command Suite

Hitachi Command Suite manages virtualized storage and server infrastructures. With usability, workflow, performance, scalability, and private cloud enablement, Hitachi Command Suite lets you build sustainable infrastructures with leading storage technologies. It helps you flexibly align with changing business requirements and maximize return on IT investments.

For this solution, the Hitachi Command Suite was used to create LDEVs and dynamic pools for the SAP Solution Manager instance and SAP ERP 6.0 Enhancement Pack 4 ABAP instance. It was used also to assign the newly created LDEVs to host groups.

For more information, see the Hitachi Command Suite Software *User Guide* shipped with the product.

Hitachi Universal Replicator

Hitachi Universal Replicator for the Hitachi Virtual Storage Platform is an advanced technology for asynchronously replicating data hosted on Virtual Storage Platform models.

Asynchronous replication maintains a remote copy of SAP data without having a performance impact on the primary data. In asynchronous replication, an I/O update operation is considered complete when it is confirmed at the primary site.

The Hitachi Virtual Storage Platform manages the process of replicating the changes to the secondary site.

Hitachi Universal Replicator software uses disk-based journaling and an optimized replication engine to reduce resource consumption and costs while increasing performance and operational resilience. The strengths of Hitachi Universal Replicator software are two key technical innovations: performance optimized, disk based journaling and a pull-style replication engine.

Following are the key characteristics of the Universal Replicator software:

- **Replication across any distance without significant negative effect on host performance**— Hitachi Universal Replicator has been successfully used in replication configurations that span thousands of miles.
- **No acknowledgement dependencies from secondary site**— Hitachi Universal Replicator replicates to a remote site without the performance impact of waiting for each individual record to be acknowledged. Instead, Hitachi Universal Replicator manages the remote relationship at a controller level. If the communication to the remote unit is disrupted or the capability of the replication circuit is exceeded, Hitachi Universal Replicator retains a replicated write data in local journals, then updates the write data when the condition is corrected.

There is a potential for some data lag between remote and primary sites, particularly at longer distances. The recovery point objective is managed with the configuration of the data communication lines. When I/O activity at the primary site exceeds the capacity of the communication channel, the data is staged and moved to the secondary site in the same order as it was written at the primary site.

Hitachi Dynamic Provisioning

On the Hitachi Virtual Storage Platform, Hitachi Dynamic Provisioning provides wide striping and thin provisioning functionalities.

Using Hitachi Dynamic Provisioning is similar to using a host-based logical volume manager (LVM), but without incurring host processing overhead. It provides one or more wide-striping pools across many RAID groups within a Hitachi Virtual Storage Platform. Each pool has one or more dynamic provisioning virtual volumes (DP-VOLs) of a user-specified logical size of up to 60TB created against it (with no initial physical space allocated).

Deploying Hitachi Dynamic Provisioning avoids the routine issue of hot spots that occur on logical devices (LDEVs). These occur within individual RAID groups when the host workload exceeds the IOPS or throughput capacity of that RAID group. This distributes the host workload across many RAID groups, which provides a smoothing effect that dramatically reduces hot spots.

Hitachi Dynamic Provisioning software has the benefit of thin provisioning. Physical space assignment from the pool to the DP-VOL happens as needed using 42MB pages, up to the logical size specified for each DP-VOL. There can be a dynamic expansion or reduction of pool capacity without disruption or downtime. An expanded pool can be rebalanced across the current and newly added RAID groups for an even striping of the data and the workload.

For more information, see the [Hitachi Dynamic Provisioning software datasheet](#) and [Hitachi Dynamic Provisioning software](#) on the Hitachi Data Systems website.

Hitachi Dynamic Link Manager Advanced

Hitachi Dynamic Link Manager Advanced combines all the capabilities of Hitachi Dynamic Link Manager and Hitachi Global Link Manager into a comprehensive multipathing solution. It includes capabilities such as the following:

- Path failover and failback
- Automatic load balancing to provide higher data availability and accessibility.

In this solution, Hitachi Dynamic Link Manager is installed on the three SAP servers.

For more information, see [Hitachi Dynamic Link Manager](#) on the Hitachi Data Systems website.

SAP Software

SAP ERP combines scalable and effective software for enterprise resource planning (ERP) with a flexible, open technology platform that can leverage and integrate SAP and non-SAP systems. SAP provides end-to-end software functionality for enterprise management.

SAP Solution Manager is a service and support platform that provides the integrated content, tools and methodologies to implement, support and monitor operations of SAP implementation. Solution Manager is installed on a separate system. Hitachi Data Systems uses the Solution Manager server for these functions:

- Configuring installation keys for the install of any SAP ABAP and Java instance
- Approving the download of any updates from SAP like service packs or enhancement packs
- Installing SAP enhancement packs onto ABAP instances
- Obtaining support from SAP

SAP graphical user interface is the client software that allows SAP users to access various functionalities in SAP applications and SAP Solution Manager.

SAP systems can be used with a variety of databases available from different vendors. The business transactions in SAP systems are processing units grouped to provide specific functions; these processing units execute changes to the database that are consistent.

Microsoft SQL Server

Microsoft SQL Server is state-of-the-art high performance database. For this solution, Hitachi Data Systems used Microsoft SQL Server 2008 R2 Enterprise Edition.

Microsoft SQL Server Management Studio is a rich, integrated administrative graphical tool client, which is designed to do all Microsoft SQL Server administrative tasks. SQL Server Management Studio is used to access, configure, manage, administrate, and develop all components of SQL Server. With its easy-to-use graphical tools, rich script editors, and command help, SQL Server Management Studio provides SQL Server access to developers and administrators of all skills.

Solution Implementation

Deploying this solution requires these following steps:

1. [Configure the storage area network.](#)
2. [Configure servers and the operating environment.](#)
3. [Configure storage for SAP ERP server at remote site.](#)
4. [Install and configure SAP ERP on server at remote site.](#)
5. [Prepare the storage system for Hitachi Universal Replicator.](#)
6. [Configure Hitachi Command Control Interface \(CCI\) for Replication.](#)
7. [Implement remote replication for SAP ERP server.](#)

There are more details about each of these steps in the following sections.

Your implementation checklist may vary, based on your environment and your requirements.

Configure the Storage Area Network

All the server blades at the local and remote site have one Fibre Channel mezzanine card installed. The mezzanine cards are connected internally to the internal Fibre Channel switch modules located in the Hitachi Compute Blade 2000 chassis.

Eight inter-switch links were connected to the Brocade 5300 switches, which were connected to eight ports of the Hitachi Virtual Storage Platform storage system.

When deploying a SAP system, isolate the storage used by the SAP Solution Manager server and the SAP ERP server at the local and remote sites from each other in a storage area network (SAN) by the use of zones and host groups. Each connection is configured with zones on fabric switches following the manufacturer's guidelines. In addition, follow these recommended practices:

- Use World Wide Port Name (WWPN) identification for all zoning configuration.
- Connect a minimum of two HBAs per server for multipath high availability.
- Disable all unused switch ports to increase security.
- Configure ports for point-to-point topology.
- Set ports to a specific speed. Do not use the auto negotiate setting.
- Use single initiator zoning.

This solution uses two *initiator* ports on the storage system at the *local site* connected to two target ports on the storage system at the remote site. Also, it uses two *initiator* ports from the storage system at the *remote site* connected to two target ports on the storage system at the local site.

Figure 2 shows the storage network configuration of the SAP instances at the local site.

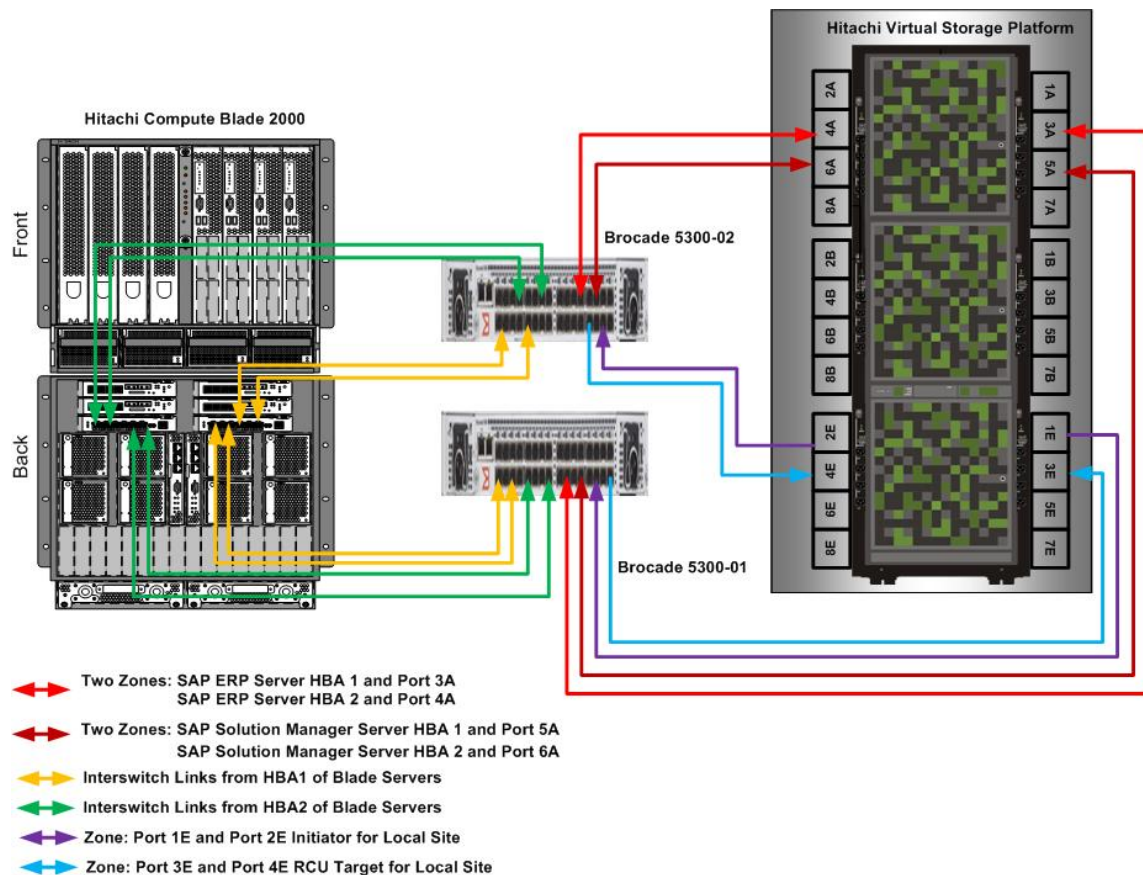


Figure 2

Figure 3 shows the storage network configuration of the SAP instances at the remote site.

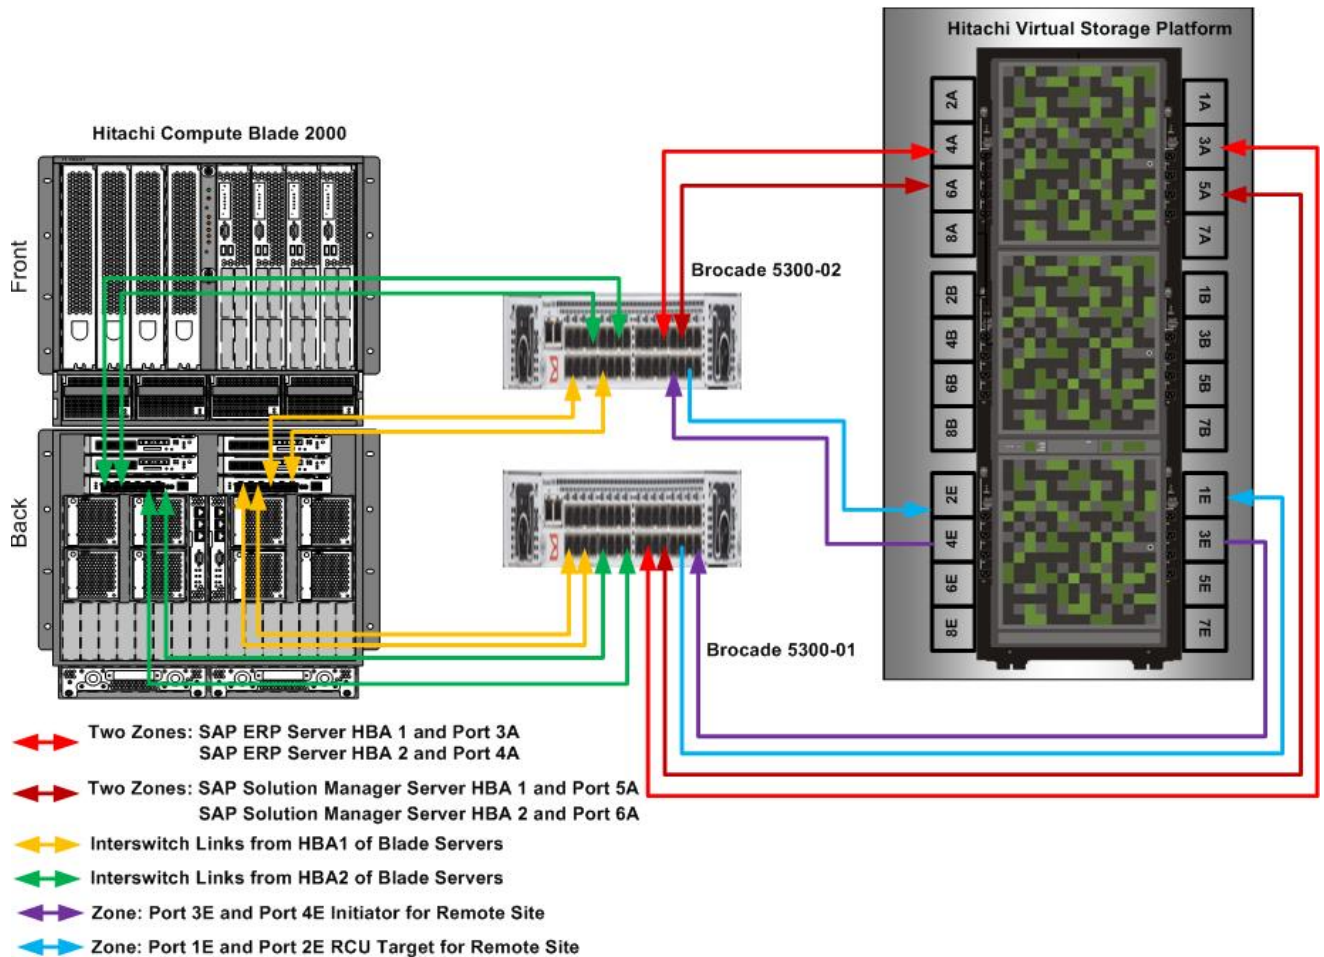


Figure 3

Configure Servers and the Operating Environment

Deploying this solution requires configuring the servers and the operating environment by implementing the following steps. Refer to these steps, as discussed in detail in [Deploying SAP on Microsoft SQL 2008 Environments Using the Hitachi Virtual Storage Platform](#).

1. Configure Management Servers.
2. Configure Hitachi Compute Blade 2000
3. Deploy the Windows Server 2008 Operating Systems
4. Deploy the Microsoft SQL Server 2008 for SAP
5. Deploy SAP Solution Manager software
6. Deploy SAP ERP ABAP Software
7. Upgrade SAP Instances

When deploying this solution, Hitachi Data Systems used the same hardware, software, and storage configuration for the SAP Solution Manager and the SAP ERP server at the local site, as described in [Deploying SAP on Microsoft SQL 2008 Environments Using the Hitachi Virtual Storage Platform Implementation Guide](#).

Deploy the SAP ERP server at the remote site following the same guidelines used for the SAP ERP server at the local site. The underlying RAID configuration, RAID type, and disk type for the storage remain the same for the SAP ERP server at the remote site. Use the same user number (UID) and the group number (GID) for the SAP software owner and Oracle users on the SAP ERP server at the local and remote sites.

Depending on the anticipated workload of the SAP ERP server at the remote site, you can use any type of supported server that meets processing needs. However, the use of highly available and scalable storage systems is a necessity for implementing SAP replication.

The SAP Solution Manager server at the remote site can be installed and configured independent of the SAP ERP server on the remote site. This paper does not include details about configuring the SAP Solution Manager server at the remote site because it is not part of the replication process.

Configure Storage for SAP ERP server at remote site.

This describes how to configure the storage for your SAP ERP server at the remote site. For this solution, Hitachi Data Systems used Hitachi Dynamic Provisioning to provision the storage for the databases.

Standard LDEVs were used for storing the transactional logs of the SAP ERP server. Dynamic Provisioning volumes (DP-VOLs) were used for all the other logical volumes of the SAP ERP server. To make sure that the failure of a disk that contains a data file does not cause the loss of the transactional logs needed to recover the data file, this solution stored the transactional logs on separate disks from the data file disks.

The dynamic provisioning pool SAP_ERP_REMOTE was created for the SAP database of the SAP ERP server at the remote site. The RAID groups used to create the dynamic provisioning pool were configured as RAID-5 (3D+1P). All drives were 300GB 10K RPM SAS drives. A total of two RAID groups were used for the dynamic provisioning pool SAP_ERP_REMOTE, giving the pool a storage capacity of 1.6TB. Six DP-VOLs were configured in the dynamic provisioning pool.

Table 1 lists storage configuration details of LDEVs created for the SAP ERP server in the SAP_ERP_REMOTE pool. All LDEVs were mapped to storage ports 3A and 4A.

Table 1. LDEV Storage Configuration for the SAP ERP Server at the Remote Site

LDEV	Host LUN	Size (GB)
00:01:50	00	200
00:01:51	01	100
00:01:52	03	125
00:01:53	04	125
00:01:54	05	125
00:01:55	06	125

One RAID-10 (2D+2P) RAID group was used to store the SAP transaction logs. All drives were 300GB 10K RPM SAS drives. The RAID group had a storage capacity of 600GB. One LDEV of 100GB size carved out of the RAID Group was used to store the Microsoft SQL Server 2008 transactional logs for the SAP ERP server at the remote site. The corresponding host LUN is 02.

Table 2 lists the file system layout details for the SAP ERP off-host server.

Table 2. File System Layout for the SAP ERP Server at the Remote Site

Host LUN	File System Mount Point	Usage
00	C: \	Host Operating System and Microsoft SQL server 2005 installation files
01	S: \usr\sap	File system for SAP binaries, SAP trans.
02	L: \<SID>LOG	SAP Transaction Log
03	M: \<SID>DATA	SAP Data 1
04		SAP Data 2
05		SAP Data 3
06		SAP Data 4

In the tested solution, dynamic disks were used for the drive M mount points on SAP ERP server at remote site for the data files and they were created as spanned disks.

Install and Configure SAP ERP on Server at Remote Site

After configuring the storage for the SAP ERP server at the remote site follow the same steps as described for the SAP ERP server in [Deploying SAP on Microsoft SQL 2008 Environments Using the Hitachi Virtual Storage Platform Implementation Guide](#) for the installation and configuration of the SAP ERP software on the server at remote site. Upgrade the SAP kernel and the SAP Microsoft Management Console to the latest release.

The SAP Microsoft Management Console (SAP MMC) provides a common framework for system management. It integrates previously separate tools in a common user interface for centralized system management. SAP has developed the SAP Systems Manager snap-in to let you monitor, start, or stop the SAP system centrally from the SAP MMC. SAP recommends that you use the SAP MMC because it simplifies system administration and provides many features.

Prepare Storage Systems for Hitachi Universal Replicator Software

Hitachi Universal Replicator Software Functionality

Hitachi Universal Replicator software is a disaster recovery solution for protecting large amounts of data that span multiple volumes. Hitachi Universal Replicator software uses group-based, update sequence consistency to enable fast and accurate recovery.

Hitachi Universal Replicator software creates duplicate volumes by copying data from the primary data volumes (P-VOL) in the storage system or main control unit (MCU) at the local site to the secondary data volumes (S-VOL) in the storage system or remote control unit (RCU) at the remote location. Hitachi Universal Replicator software also makes sure that any updates on the primary data volumes are applied to the secondary data volumes.

Hitachi Universal Replicator software performs the following operations to copy data from the primary data volumes to the secondary data volumes and maintain replication:

- **Base journal copy**—Also known as the initial copy, this operation synchronizes the data in the P-VOL and the S-VOL. This operation is independent from host I/O. It is performed when a data volume pair is created or a suspended pair is resynchronized.
- **Journal obtain**—Stores the write data as a journal data in the primary journal volume (M-JNL) in the MCU with every update of the P-VOL, according to the write instructions from the host. The host assigns write-sequence numbers to the data sent to the master journal volume. Write-sequence numbers and the metadata attached to the journal data insure consistency with the data in the P-VOL.
- **Journal copy**—Copies data in the M-JNL in the MCU to the secondary journal volume (R-JNL) at the remote site. During the journal copy operation, data is pulled from the MCU journal cache to the RCU journal cache and then to the R-JNL (data might be read from the M-JNL if the data was destaged to the M-JNL). The RCU issues a read journal command to the MCU to request the transfer of the journal data that is stored in the primary journal volumes after the pair create or pair resync operation from the MCU. After the journal copy operation is complete, the MCU discards the journal data.
- **Journal restore**—Writes data from the R-JNL to the S-VOLs volumes by maintaining the write order fidelity. After the data in the R-JNL is written to the S-VOL, the RCU discards the journal data.

Figure 4 describes Hitachi Universal Replicator operations.

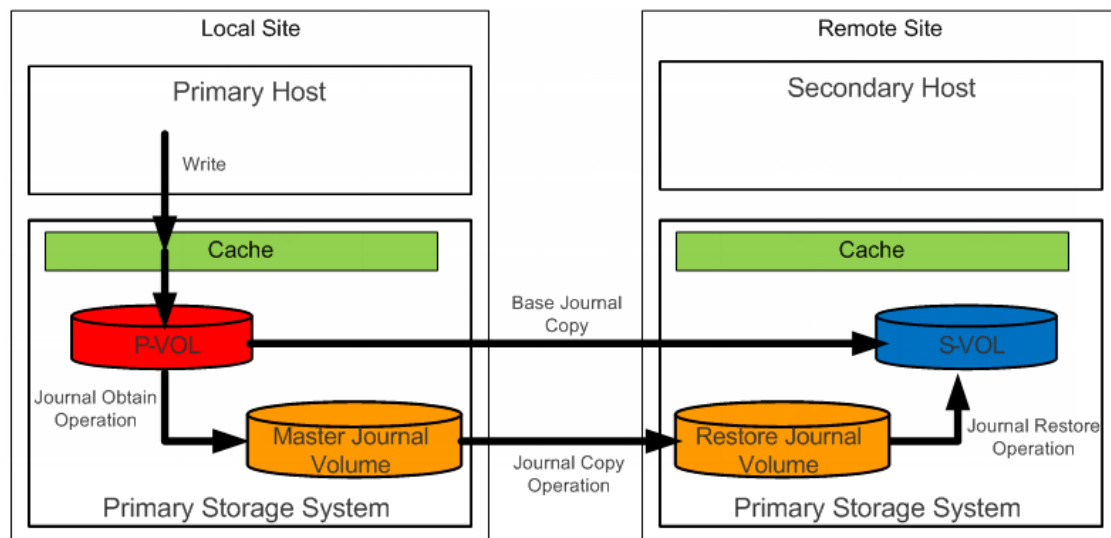


Figure 4

To determine the required number of initiator and RCU target ports on a Hitachi Virtual Storage Platform refer to the “*Hitachi Universal Replicator User Guide*” for the Hitachi Virtual Storage Platform. To enable Hitachi Universal Replicator software on your storage systems follow these steps on both the MCU and RCU storage systems using Hitachi Storage Navigator software:

1. [Configure port attributes.](#)
2. [Add a disk controller \(DKC\) definition.](#)
3. [Create journal groups.](#)

Configuring Port Attributes

You must have at least one initiator and one RCU target. This solution uses two initiators and two RCU targets at each site. Configure port attributes of the MCU and RCU storage systems using Hitachi Storage Navigator.

To convert MCU target ports to initiators or RCU targets, follow these steps on the MCU.

1. From the **Actions** menu, point to **Remote Copy**, then point to **Universal Replicator**, and click **DKC Operation**.

The **Universal Replicator DKC Operation** window opens.

2. Click the **Port** option.

The **Subsystem** navigation tree displays in the left area.

3. Convert one or more target ports to be an initiator port.
 - (1) In the right area, right-click the port to be converted to an initiator (for example, 2E), and click **Initiator** from the list.

The port changes to an initiator.

- (2) Click **Apply**.

A dialog box displays asking **Do you want to apply?**

- (3) Click **OK**.

A dialog box displays asking **Are you sure you want to exchange the specified ports? Please confirm the associated channel paths are off-line.**

- (4) Click **OK**.

A dialog box displays asking **Exchange the channel cable connections after confirmation of the port type.**

- (5) Click **OK**.

- (6) **The requested operation is complete** prompt displays in a pop-up window

Click **OK**.

Repeat these steps to convert additional targets to initiators.

4. Convert one or more target ports to be an RCU target.
 - (1) In the right area, right-click the port to be converted to an RCU target (for example, 4E), and click **RCU Target** from the list.

The port changes to an initiator.

- (2) Click **Apply**.

A dialog box displays asking **Do you want to apply?**

- (3) Click **OK**.

A dialog box displays asking Are you sure you want to exchange the specified ports? Please confirm the associated channel paths are off-line.

- (4) Click **OK**.

A dialog box displays asking Exchange the channel cable connections after confirmation of the port type.

- (5) Click **OK**.

- (6) The requested operation is complete prompt displays in a pop-up window

Click **OK**.

Repeat these steps to convert additional targets to initiators.

5. Repeat the previous steps to configure initiator and target ports on the RCU storage system.

Adding a DKC Definition

To add a DKC definition on the MCU and RCU storage systems, follow these steps using Hitachi Storage Navigator software on the MCU storage system:

1. From the **Actions** menu, point to **Remote Copy** and then **Universal Replicator**, and then click **DKC Operation**.

The **Universal Replicator DKC Operation** window opens.

2. Select the **DKC** option.

The **LDKC** navigation tree displays in the left area.

3. Click **LDKC#00** on the navigation tree.

4. Right-click a blank space in the left area.

A menu displays.

5. From the menu, click **DKC Operation** and then **Add DKC**.

The **ADD DKC** dialog box opens.

6. Provide the following information in the **ADD DKC** dialog box:

- (1) Type the serial number for the remote storage system in the **S/N** field.

- (2) Choose the LDKC of the remote storage system from the **LDKC** drop-down menu.

- (3) Choose **6 (VSP)** from the **Controller ID** drop-down menu.

- (4) Choose the local storage system ports and its corresponding pair-port in **M-R path** pane.

- (5) Click **Option**.

The **DKC Option** dialog box displays.

- (6) Enter values in the **Minimum Paths** and **RIO MIH Times (sec)** fields.

(7) Click **Set**.

The **DKC Option** dialog box closes.

7. Click the **Apply** button.

A dialog box displays asking **Do you want to apply?**

8. Click **OK**.

A dialog box displays saying **The requested operation is complete.**

9. Click **OK**.

Repeat these steps on the RCU storage system to add a DKC there.

Creating Journal Groups

To create journal groups using Hitachi Storage Navigator, follow these steps on the MCU storage system:

1. From the **Actions** menu, point to **Remote Copy** and then **Universal Replicator**, and then click **DKC Operation**.

The **Universal Replicator Journal Operation** dialog box opens.

2. In the **Journals** navigation tree, click to expand **Free**.

3. Select a free journal group from the list to highlight it.

Details of the selected free journal group display in the right pane.

4. Right-click the highlighted journal group in the right area, and click **Edit JNL VOLs** from the menu.

The **Edit JNL Volumes** dialog box opens.

5. Click the **CU** option.

6. From the menu to the right of the **CU** open, click a CU.

The **Free Volumes** pane populates.

7. Select one or more volume in the **Free Volumes** area.

The **Add** button becomes available.

8. Click **Add**.

The **JNL Volumes** area populates.

9. Click **Set** and then Click **Apply**.

A dialog box displays saying **Do you want to apply?**

10. Click **Yes**.

A dialog box displays saying **The requested operation is complete.**

11. Click **OK**.

Repeat these steps on the RCU storage system to add a journal group on the RCU.

Configure Hitachi Command Control Interface (CCI) for Replication

Although Hitachi Universal Replicator residing on the MCU and RCU storage systems maintains the replication of data between the P-VOLs on MCU and the S-VOLs on RCU, it is Hitachi Command Control Interface that allows you to manage, monitor, and control the Hitachi Universal Replicator process.

The configuration of Hitachi Command Control Interface involves the following steps:

1. [Install Hitachi Command Control Interface.](#)
2. [Configure command devices.](#)
3. [Configure replication using Hitachi Command Control Interface.](#)

Install Hitachi Command Control Interface

To install Hitachi Command Control Interface on the SAP ERP server on the local and remote site, follow these steps:

1. Insert the installation media for the Hitachi Command Control Interface into an I/O device (for example, CD-ROM).
2. Run **Setup.exe** on the installation media. Follow the on-screen instructions to complete the installation.
3. Verify installation of the latest version of Hitachi Command Control Interface by typing the following command at a command prompt from the Open Remote Copy manager installation directory (C: \HORCM\etc):

```
rai dqry -h
```

The output looks something like the following:

```
Model   : RAID- Manager/WindowsNT  
Ver&Rev: 01-24-03/06  
Usage   : rai dqry [options] for HORCM
```

Configure Command Devices

A command device is a dedicated logical volume on the storage system that functions as the interface to the storage system from the host. The command device accepts commands from the host that are executed on the storage system.

Before doing these steps, create a 1GB LDEV for the command device from the RAID group used to store the Microsoft SQL Server 2008 transactional logs. Associate this LDEV to the host group for the host connection using Hitachi Storage Navigator.

Convert the LDEV to a command device in Hitachi Storage Navigator by following these steps:

1. Insert the installation media for the Hitachi Command Control Interface into an I/O device (for example, CD-ROM).
2. Click **Logical Devices** in the **Explorer** area.
3. In the right area on the **LDEVs** tab, select the LDEV to be converted to command device.

4. Click **More Actions Options** and then click **Edit Command Device**.

The Edit Command Devices dialog box opens.

5. In the **Edit Command Devices** dialog box, click the **Enable** option.
6. Click **Finish**.
7. Click **Apply**.

Repeat these steps to create another LDEV from the RAID group used to store the Microsoft SQL Server 2008 transactional logs of the SAP ERP server at remote site. Then, map the LDEV to the host group and convert it to a command device so that Hitachi Command Control Interface can use it to configure replication between the SAP ERP server at both the local site and remote site.

The command device must be defined and configured as a raw device with no file system and no mount operation.

Configure Replication Using Hitachi Command Control Interface

Hitachi Command Control Interface has these components residing on the following:

- **Storage system**—Command devices and Hitachi Universal Replicator volumes (P-VOLs and S-VOLS)
- **Server**—Open Remote Copy Manager (commonly called HORCM) configuration definition files (for example, `horcm1.conf`), and Hitachi Command Control Interface commands

The Open Remote Copy Manager operational environment operates as a daemon process on the host server. When activated, the Open Remote Copy Manager refers to the configuration definition files. The Open Remote Copy Manager instance communicates with the storage sub-system and remote servers.

Two instances of Open Remote Copy Manager are required for Hitachi Universal Replicator to be operational. One instance manages the P-VOLs, while the other instance manages the S-VOLs. The Open Remote Copy Manager configuration file defines the communication path and the logical units to be controlled. Each instance has its own configuration file. The content of the `horcm0.conf` and the `horcm1.conf` files used for this solution are available in "[Appendix A — Hitachi Universal Replicator Configuration Files](#)".

Modify the `services` file in the `C:\Windows\System32\drivers/etc\` folder to register the port name and number for each Open Remote Copy Manager instance on each server. The port name entries for Open Remote Copy Manager (HORCM) in the `services` file must be the same on all the servers. For example, if the service number for port name `horcm0` is `11000/udp`, the service number for port name `horcm1` must be `11001/udp` in the `services` file on SAP ERP server at the local and the remote site.

Figure 5 shows a two-server, two-HORCM instance configuration used for Command Control Interface for this solution.

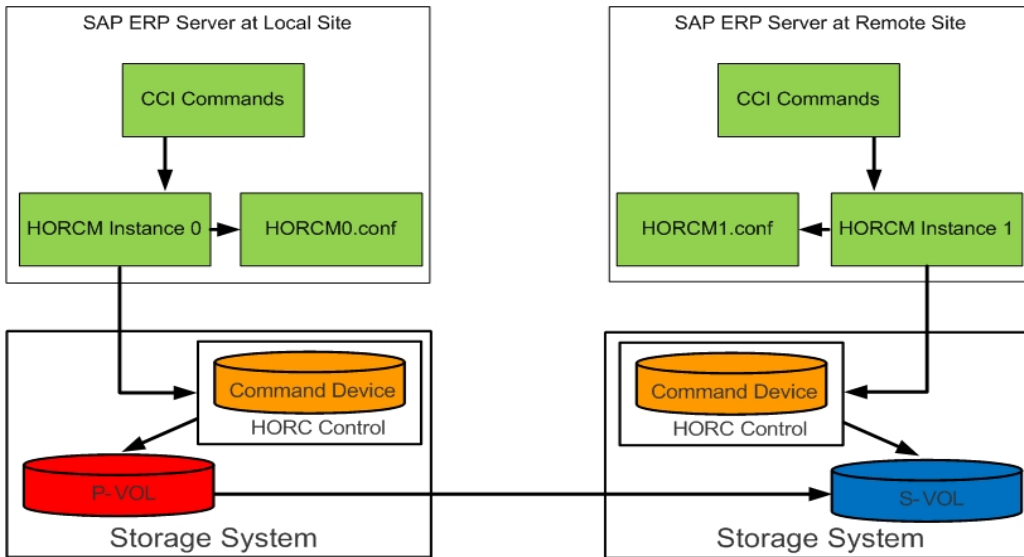


Figure 5

For more information about setting up the Open Remote Copy Manager instances and Open Remote Copy Manager configuration file, see the “Hitachi Command Control Interface (CCI) Users and Reference Guide” that accompanies the Hitachi Command Control Interface software.

CCI commands are executed as windows administrator user. Verify that that the environmental variables HORCM_CONF, HORCMINST and HORCC_MRCF are set for the instance on each server and then start the HORCM instance on each server. To set the above mentioned environment variables and to start up the HORCM instance follow these steps:

Note— Set the environmental variables for each instance. In the following instructions, [*] is substituted for the instance number in the instructions. For example, to set HORCM_CONF for instance 2, type SET HORCM_CONF=C:\HORCM\etc\horcm*.conf to set the environment variable.

1. Open a Command Prompt window.
2. At the command prompt, type the following to set HORCM_CONF:
`SET HORCM_CONF=C:\HORCM\etc\horcm*.conf`
3. At the command prompt, type the following to set HORCMINST(instance number to start):
`SET HORCMINST=*`
4. At the command prompt, type the following to start Hitachi Open Remote Copy Manager instance:
`cd C:\HORCM\etc`
`horcmstart`

After starting the Open Remote Copy Manager instance on all the servers, execute CCI commands to manage, monitor, and control replication between the servers. Table 3 lists CCI commands for pair operations that can be executed in a Command Prompt window.

Table 3. CCI Commands for PAIR Operatin

<i>Command Name</i>	<i>Description</i>	<i>Notes</i>
<code>paircreate</code>	The paircreate command starts the Hitachi Shadow Image initial copy operation, which overwrites all data on the secondary/target volume. Generates a new volume pair from two unpaired volume.	Status of P-VOL and S-VOL changes from SMPL to PAIR.
<code>pairsplit</code>	Splits pairs and create a point-in-time backup in the S-VOL by stopping the update of the secondary volume. The pair relationship of the P-VOL is maintained, and the differential information is also maintained.	Status of P-VOL changes from PAIR or COPY to PSUS. Status of S-VOL changes from PAIR or COPY to SSUS.
<code>pairsplit -S</code>	Discards the pair's differential information and deletes the pair.	Status of P-VOL changes to SMPL.
<code>pairdisplay</code>	Checks pair status and progress rate.	Use -l option in event of a failure when one of the sites is unavailable
<code>horctakeover</code>	Takes over operations from the local site to the remote site. It is also used to reverse the copy direction and re-establish the pair.	In the event of a local site failure, the status of S-VOL changes from PAIR to SSWS In the event of a planned failover, P-VOL becomes S-VOL and S-VOL becomes P-VOL
<code>paircreate</code>	Creates a new pair.	Status of P-VOL and S-VOL changes from SMPL to PAIR.
<code>pairsplit</code>	Splits pairs and create a point-in-time backup in the S-VOL by stopping the update of the secondary volume. The pair relationship of the P-VOL is maintained, and the differential information is also maintained.	Status of P-VOL changes from PAIR or COPY to PSUS. Status of S-VOL changes from PAIR or COPY to SSUS.

Implement Remote Replication for SAP ERP Server

The SAP ERP system deployment consists of the following:

- SAP binaries
- Microsoft SQL Server installation files
- Data in the database or databases

SAP ERP binary files and Microsoft SQL Server installation files are deployed on separate LUNs from the LUNs containing database data. Because the solution described in this white paper installs and configures SAP ERP server at the remote site, the Open Remote Copy Manager configuration files are defined in such a way that the Hitachi Universal Replicator creates pairs of the volumes containing database data only. Database data includes data files and the log files.

For remote replication of the data on the SAP ERP server, initiate the base journal process by using the `paircreate` command on the SAP ERP server at the local site:

```
paircreate -g <group name> -vl -f <fence level> -jp -js <journal group ID>
```

For example:

```
paircreate -g SAP-HURG01 -vl -f async -jp 00 -js 00
```

When the `paircreate` command is executed, the initial copy is executed. The primary storage system copies all the data in sequence from the P-VOL directly to the S-VOL. During the initial copy process the status of the P-VOL and S-VOL is COPY. On completion of the initial copy process, the status of the P-VOL and S-VOL changes to PAIR. Any subsequent updates to the P-VOL are transferred to the S-VOL following the process described in the “Hitachi Universal Replicator Software Functionality” section.

Recovering the SAP ERP Server During a Planned Failover

You might plan a failover of an SAP ERP server from a local site to a remote site for various reasons such as site migration, temporary failover to remote site for maintenance of local site, and so on. In a planned failover the applications at the local site are shutdown gracefully, Hitachi Universal Replicator software replicates any pending data that needs to be transferred to the remote site and applications are started at the remote site before servers and storage are shutdown on the local site.

To recover the SAP ERP server on the remote site during a planned failover, follow these steps:

1. On the SAP ERP server at the local site, open the SAP Microsoft Management Console (SAP MMC), right click the SAP ERP off-host SID and click **Stop**.

The **SAP System Shutdown** dialog box opens.

2. Choose the **Soft** option and click **OK**.

The **Webservice Authentication** dialog box opens.

3. Type the user ID and password for the SAP SID administrator and click **OK**.
4. Using the Microsoft Windows disk management system utility, unmount those file systems containing the off-host server's database data and log files.

- In the **Disk Management** area, right-click the disks containing the database and transactional log files and then click **Offline**.

5. On the remote site, execute the `horctakeover` command:

```
# horctakeover -g SAP-HURG01 -t 60 -S
```

The takeover switch function evaluates the status of P-VOLs and S-VOLs. Because the local volume is the S-VOL and the remote volume is the P-VOL and the status is PAIR, a SWAPTAKEOVER function is selected and the S-VOL is suspended to allow swapping of the P-VOL. Resync for the S-VOL is executed so that the S-VOL becomes the P-VOL. The new P-VOL is resynchronized with the new S-VOL.

After the successful takeover, perform the following procedure to make sure that the SAP ERP server at the remote site is functional:

1. Mount the file system containing database data on the SAP ERP server at the remote site using the Windows Disk Management tool.
 - In the **Disk Management** area, right-click the disks containing the transactional log file and click **Online**.
The data files are carved out of dynamic disks. Mount these dynamic disks.
 - Right-click the dynamic disks containing the data files, and then click **Import Foreign Disks**.

When the file system is mounted on the ERP server at the remote site, the system provides a new drive letter for the disk containing transactional log files as they are configured as basic disks.

2. Change the drive letter back to the original drive letter used for the disks containing the transactional logs.
 - Right-click the drive, click **Change Drive Letter and Paths**, and then type the correct drive letter.
3. Log on to the Microsoft SQL Server Management Studio and restart the MSSQLSERVER service on the SAP ERP server at the remote site.
4. To assign the necessary authorizations for the SQL server logins, after the database copy between two SAP systems with the same SID execute the script file user_restore.sql. Refer to “Appendix B—Assigning Authorizations for SQL Server Logon After Replication” for details on the script file user_restore.sql .
5. Assign the necessary SQL server authorizations for the copied database:
 - Click **New Query** to execute the script user_restore.sql .
6. To start the SAP instance from SAP MMC, right-click the SAP off-host SID and then click **Start**.
The **SAP System Start** dialog box opens.
7. Click **OK**

The SAP ERP system at the remote site is now available for business processing.

Recovering SAP ERP Server After a Local Site Disaster

In a disaster scenario where the local site is completely unavailable and the links between the local site and remote site are broken, it is possible that some of the data from the MCU might not be applied at the RCU. In this scenario, an incomplete recovery of the database must be performed at the remote site leading to loss of some data.

To recover SAP ERP server on the remote site after a disaster at the local site, follow these steps:

1. On the remote site execute horctakeover command as a windows administrator user:

```
# horctakeover -g SAP-HURG01 -t 60
```

The takeover switch function determines that the command was executed at the S-VOL site and tries to determine the status of P-VOL. Because the status of P-VOL is unknown, the takeover function executes an S-VOL takeover and changes the status of S-VOL from PAIR to SSWS with an assumption that local site cannot be accessed. The S-VOL with SSWS status can be mounted for host I/O operations.

2. Mount the file system containing database data on the SAP ERP server at the remote site using the Windows Disk Management tool.
 - In the **Disk Management** area, right-click the disks containing the transactional log file and click **Online**.

The data files are carved out of dynamic disks. Mount these dynamic disks.

- Right-click the dynamic disks containing the data files, and then click **Import Foreign Disks**.

When the file system is mounted on the ERP server at the remote site, the system provides a new drive letter for the disk containing transactional log files as they are configured as basic disks.

3. Change the drive letter back to the original drive letter used for the disks containing the transactional logs.
 - Right-click the drive, click **Change Drive Letter and Paths**, and then type the correct drive letter.
4. Log on to the Microsoft SQL Server Management Studio and restart the MSSQLSERVER service on the SAP ERP server at the remote site.
5. To assign the necessary authorizations for the SQL server logins, after the database copy between two SAP systems with the same SID execute the script file `user_restore.sql`. Refer to “Appendix B—Assigning Authorizations for SQL Server Logon After Replication” for details on the script file `user_restore. sql` .
6. Assign the necessary SQL server authorizations for the copied database:
 - Click **New Query** to execute the script `user_restore. sql` .
7. To start the SAP instance from SAP MMC, right-click the SAP off-host SID and then click **Start**.
The **SAP System Start** dialog box opens.
8. Click **OK**

The SAP ERP system on the remote is now available on the remote site for business processing.

Appendix A—Hitachi Universal Replicator Configuration Files

Hitachi Universal Replicator configuration files consist of following sections:

- **HORCM_MON**—Contains information needed to monitor a HORCM instance such as IP address, HORCM instance or service, pooling interval for monitoring paired volumes, and timeout period for communication with remote server.
- **HORCM_CMD**—Contains device path information about the command device.
- **HORCM_LDEV**—Defines the storage sub-system device address for the paired logical volume names.
- **HORCM_INST**—Network address of the remote server.

Copy the following `horcm0.conf` configuration file to `C:\HORCM\etc` folder on ERP server at the local site and `horcm1.conf` configuration file to `C:\HORCM\etc` folder on ERP server at the remote site.

HORCM Configuration File for SAP ERP Server at Local Site

```
#FileName: horcm0.conf
#/****** For HORCM_MON *****/
HORCM_MON
#ip_address          service      poll (10ms)      timeout (10ms)
172.17.171.213      horcm0        1000              3000
#/****** For HORCM_CMD *****/
HORCM_CMD
#dev_name            dev_name      dev_name
\\.\Physical Drive7
#/****** For HORCM_LDEV *****/
HORCM_LDEV
#dev_group           dev_name      Serial#          CU: LDEV(LDEV#)  MU#
SAP-HURG01          SAP_Log_1     53101           01:15            0
SAP-HURG01          SAP_Data_1    53101           01:16            0
SAP-HURG01          SAP_Data_2    53101           01:17            0
SAP-HURG01          SAP_Data_3    53101           01:18            0
SAP-HURG01          SAP_Data_4    53101           01:19            0
#/****** For HORCM_INST *****/
HORCM_INST
#dev_group           ip_address    service
SAP-HURG01          172.17.171.214 horcm1
```

HORCM Configuration File for SAP ERP Server at Remote Site

```
#FileName: horcm1.conf
#/****** For HORCM_MON *****/
HORCM_MON
#ip_address          service      poll (10ms)      timeout (10ms)
172.17.171.214      horcm1        1000              3000
#/****** For HORCM_CMD *****/
HORCM_CMD
#dev_name            dev_name      dev_name
\\.\Physical Drive7
#/****** For HORCM_LDEV *****/
HORCM_LDEV
#dev_group           dev_name      Serial#          CU: LDEV(LDEV#)  MU#
SAP-HURG01          SAP_Log_1     53119           01:51            0
SAP-HURG01          SAP_Data_1    53119           01:52            0
SAP-HURG01          SAP_Data_2    53119           01:53            0
SAP-HURG01          SAP_Data_3    53119           01:54            0
SAP-HURG01          SAP_Data_4    53119           01:55            0
#/****** For HORCM_INST *****/
HORCM_INST
#dev_group           ip_address    service
SAP-HURG01          172.17.171.213 horcm0
```

Appendix B—Assigning Authorizations for SQL Server Logon After Replication

To assign the necessary authorizations for the SQL server logon for the copied database after the replication, refer to the SAP Note 551915—R/3 won't start after database restore or database copy

In this solution both the SAP ERP at the local and remote site has the same SAP SID and both are MCOD (multi-schema) systems. All the tables in such systems belong to user 'sid' (the system ID in lower case).

According to Case 2, MCOD to MCOD in SAP Note 551915, to assign the necessary SQL server authorizations after the replication, do the following:

- Replicate the SAP system with the same SID: run `user_Restore.sql`
- Follow the steps provided in Case 2, MCOD to MCOD, of this SAP Note to execute these scripts successfully. Check the following while executing the scripts:
 1. Verify that there is a local Windows account `SAPMssXPUser` that is member of the local group "Administrators".
 2. Check the registry key setting `HKEY_LOCAL_MACHINE\SOFTWARE\SAP\<SID>\AdmUser` to verify it contains a string with the value `DOMAIN\sidadm`. If this key does not exist, create it.
 3. Open the script and replace variables described in the header of the script.
 4. Execute the script.
 5. Check that the script ran correctly.

After executing the scripts successfully follow up the remaining steps in the section “Recovering the SAP Server During a Planned Failover” and “Recovering SAP ERP Server after a Local Site Disaster” to continue with the post replication steps.

Refer to these SAP Notes on the procedure for copying a SQL server database and restoring it:

- SAP Note 151603 – Copying an SQL Server database
- SAP Note 193816 – Restore with SQL Server

 **Hitachi Data Systems Corporation**

Hitachi is a registered trademark of Hitachi, Ltd., in the United States and other countries. Hitachi Data Systems is a registered trademark and service mark of Hitachi, Ltd., in the United States and other countries. All other trademarks, service marks and company names mentioned in this document are properties of their respective owners.

Notice: This document is for informational purposes only, and does not set forth any warranty, expressed or implied, concerning any equipment or service offered or to be offered by Hitachi Data Systems Corporation

© Hitachi Data Systems Corporation 2010. All Rights Reserved. AS-097-00 July 2011

Corporate Headquarters

750 Central Expressway,
Santa Clara, California 95050-2627 USA
www.hds.com

Regional Contact Information

Americas: +1 408 970 1000 or info@hds.com
Europe, Middle East and Africa: +44 (0) 1753 618000 or info.emea@hds.com
Asia Pacific: +852 3189 7900 or hds.marketing.apac@hds.com